



QILAHK

SECURITY YOU CAN PROVE

ADVANCED RESEARCH WHITE PAPER

QTA-R01 / 01 OF 05

Evidence Aware Trust Computation

A research model for uncertainty dependency and accountable decisions

Version 1.1 | October 12, 2026

QILAHK Trust Authority

LOCK / KEY / PROOF

qilahk.com | blog.qilahk.com

Table of contents

Evidence Aware Trust Computation

Abstract	3
Problem and research boundary	4
Why evidence aggregation is difficult	4
A claim and evidence contract	4
Position relative to existing work	4
Dependency aware appraisal mechanism	5
Construct the graph	5
Compose claims conservatively	5
Illustrative case and adversarial analysis	6
An endpoint with duplicated support	6
Attacker capabilities and trust assumptions	6
Failure modes requiring investigation	6
Evaluation and falsifiable hypotheses	7
Build a controlled evidence corpus	7
Test two primary hypotheses	7
Research roadmap and references	8
Prototype sequence	8
Open research questions	8
References	8

Abstract

Evidence Aware Trust Computation

A trust decision needs a representation of what the evidence establishes, what it contradicts, and what remains unknown. Compressing these distinctions into a single score can conceal missing proof or count several copies of one observation as independent support. This paper proposes a QILAHK research model that evaluates scoped claims through an evidence dependency graph and produces an explanation that an authorized reviewer can reconstruct.

The proposed contribution is a compositional evaluation contract: every result retains its subject, criteria version, supporting and opposing evidence, dependency lineage, and time boundary. The contract separates evidence classification from the authority decision and from resource authorization. Its value must be tested against simpler Boolean and weighted-score baselines using correlated, missing, stale, and conflicting evidence.

The central research question is whether dependency-aware explanations reduce unjustified reliance without making reviews impractical. No performance results or security proofs are asserted. The design provides explicit invariants, illustrative cases, and an evaluation protocol for a future prototype.

Intended audience

Security researchers, trust architects, and authority reviewers

Research scope

This white paper presents proposed QILAHK research mechanisms, illustrative scenarios, and evaluation protocols. Its hypotheses are intended for prototype testing and independent review.

QTA-R01 | Version 1.1 | October 12, 2026

Problem and research boundary

Why evidence aggregation is difficult

An endpoint may appear compliant in an inventory system, a dashboard, and a compliance report. These three artifacts may all derive from the same collector. Counting them as three independent observations inflates the apparent support for a claim. Conversely, one contrary observation may concern a different device revision and should not be treated as a direct contradiction until identity and scope are reconciled.

The research boundary is the appraisal of evidence for an identified claim. It excludes deciding business risk appetite, granting operational access, and asserting properties that the collection methods cannot observe. A reviewer may accept a documented exception, but that act does not transform missing technical proof into an established property.

A claim and evidence contract

Represent a claim using a stable subject identifier, a property, a context, an applicable interval, and a criteria version. An evidence item carries an artifact reference, collector identity, observation time, collection method, scope, integrity information, and derivation links. Identity resolution is a required precondition for comparing evidence across sources.

The proposed classification uses two flags: qualified support exists and qualified opposition exists. Neither flag produces unestablished; support alone produces supported; opposition alone produces refuted; both produce conflicted. These are appraisal labels under selected criteria, not probabilities or guarantees. A collection failure does not set the opposition flag unless the criteria explicitly define that failure as relevant negative evidence.

Position relative to existing work

RFC 9334 distinguishes the Attester, Verifier, and Relying Party and separates Evidence from Attestation Results [1]. This proposal borrows that role separation while adding a claim explanation contract for heterogeneous evidence. It does not claim that its labels are an IETF standard or that a deployment conforms to RATS without a separate interoperability assessment.

Dependency aware appraisal mechanism

Construct the graph

Build directed derivation edges from original observations to normalized values and reports. Group evidence by known common origins, such as the same measurement event or collector execution. Preserve the original artifacts and transformations so that a reviewer can distinguish a new observation from a new presentation of an old observation.

For each claim, apply deterministic eligibility gates before classification. The gates check subject binding, permitted collection method, integrity verification where required, freshness, scope coverage, and criteria compatibility. An ineligible item can remain visible as context but must not silently contribute qualified support. Unknown lineage should be recorded as unknown rather than assumed independent.

Compose claims conservatively

For an all-required-properties profile, every mandatory claim must be supported and no mandatory claim may be conflicted. A single refuted claim prevents the profile from being established. An unestablished claim leaves the profile unestablished. Other composition rules, including alternatives or thresholds, require explicit policy definitions and a recorded version.

The graph produces a review packet rather than a universal trust score. The packet identifies the outcome of each required claim, eligible evidence references, excluded artifacts and reasons, unresolved contradictions, and the criteria used. An authority workflow then records the disposition, rationale, conditions, and authorized decision maker.

Proposed appraisal states

Support	Opposition	State	Review implication
No	No	Unestablished	Collect proof or narrow the claim
Yes	No	Supported	Review limits and authority
No	Yes	Refuted	Investigate the observed failure
Yes	Yes	Conflicted	Reconcile qualified disagreement

An essential invariant is that changing the display format of an existing observation does not increase its independence. A second invariant is that the explanation retains the evidence and policy versions actually used. A third is that an exception is represented as an authority action, with its own scope and duration, rather than as an alteration of the underlying evidence classification.

Illustrative case and adversarial analysis

An endpoint with duplicated support

Consider an illustrative device D17 with a claim about an approved boot configuration. Collector event E1 reports the expected state. Two downstream systems import E1 and display a passing result. The graph records three artifacts but one observation lineage. It cannot infer independent confirmation merely from the number of visible reports.

A later authenticated observation E2 reports a different state. If E1 is outside the required interval, it becomes historical context and E2 may refute the current claim. If both are eligible within the claim interval and their methods are comparable, the claim is conflicted until the discrepancy is resolved. If E2 concerns a replacement device, identity reconciliation is required before either interpretation is justified.

Attacker capabilities and trust assumptions

The threat model permits an attacker to duplicate artifacts, replay old observations, submit incomplete provenance, compromise a non-root collector, and influence a dashboard transformation. It assumes the prototype can identify some protected reference artifacts and validate the cryptographic mechanisms chosen for them. A fully compromised root collector can produce authentic but false statements; the appraisal layer cannot repair that fact without independent evidence.

Qualification gates must therefore expose their assumptions. A verified signature establishes a relationship to a key under a validation policy; it does not prove the physical property described by the signer. The evaluator should retain collector assurance as a separate attribute and should avoid treating all signed artifacts as equally credible.

Failure modes requiring investigation

An attacker may split one lineage into several apparent collectors or exploit aliases to make a contradiction look irrelevant. A legitimate collector upgrade may create apparent conflict through changed semantics. The prototype must test both adversarial and accidental errors, preserve unresolved lineage, and support reviewer correction without rewriting the historical explanation.

The mechanism remains limited by the completeness of its graph. Missing dependency edges can still lead to overstated independence. A conservative mode can require independent-source claims to remain unestablished when their independence cannot be demonstrated, but the operational cost of that rule must be measured.

Evaluation and falsifiable hypotheses

Build a controlled evidence corpus

Generate claim sets with known subject identities and ground-truth property states. Vary duplication, common-source correlation, stale observations, contradictory methods, identity ambiguity, and collector compromise. Freeze criteria and eligibility rules before each experiment. Use separate scenario families for development and held-out evaluation so rule tuning does not optimize against the test set.

Compare a Boolean all-pass baseline, a weighted-score baseline with declared weights, and the dependency-aware proposal. Feed each baseline the same raw observations and document any additional metadata unavailable to it. Include an ablation that removes dependency edges to determine whether gains come from lineage handling rather than a stricter default disposition.

Test two primary hypotheses

H1 proposes that the dependency-aware method produces fewer established profiles when mandatory claims are unsupported, particularly under duplication and missing lineage. H2 proposes that reviewers reconstruct the decision basis more accurately from the graph explanation than from a score alone. Reject either hypothesis if the observed benefit is absent or is explained entirely by more frequent abstention.

Pair false establishment with abstention and review burden. A system that never establishes a claim trivially avoids false establishment but is not useful. Report the tradeoff across declared operating points, along with uncertainty intervals and scenario-level results. Measure disagreement between reviewers and retain the adjudication method used to create reference labels.

Evaluation measures and denominators

Measure	Definition
False establishment	Established profiles with a failed mandatory ground-truth condition divided by all established profiles
Coverage	Supported or refuted profiles divided by all evaluated profiles
Reconstruction accuracy	Correctly identified evidence and criteria links divided by required links in the review task
Review cost	Median and upper-percentile review time plus manual corrections per case

Report processing latency and graph size separately from decision quality. Preregister failure criteria, publish the scenario generator where feasible, and exclude any claim of production assurance until deployment-specific collectors, identity binding, and authority controls have been independently assessed.

Research roadmap and references

Prototype sequence

First implement the claim schema and deterministic eligibility gates. Then introduce provenance edges and source groups, followed by explanation export and human review tasks. Only after the appraisal invariants pass should the prototype connect to an authority workflow or publish a consumable trust representation.

Model checking or property-based testing should target duplication invariance, missing-evidence preservation, deterministic replay under pinned inputs, and explicit conflict handling. Store reviewer corrections as additional events so the reference label and the original machine appraisal remain distinguishable.

Open research questions

How can independence be represented when only partial lineage is available? Which graph details most improve reviewer accuracy? How should conflicting collection methods be reconciled without silently privileging one source? Can explanations be shared across organizations while protecting sensitive evidence and preserving enough context for meaningful review?

These questions define a research agenda rather than a completed assurance mechanism. The proposed result is a more inspectable basis for trust decisions. Its adoption should depend on demonstrated improvements in decision quality, reviewability, and operational cost under the intended threat model.

References

[1] **IETF RFC 9334 Remote ATtestation procedureS Architecture**

Role separation and appraisal terminology.

<https://www.rfc-editor.org/rfc/rfc9334>

[2] **W3C PROV DM The PROV Data Model**

Entity activity and agent relationships for provenance.

<https://www.w3.org/TR/prov-dm/>

[3] **QILAHK Trust Authority Journal**

QILAHK architectural framing and claim scope.

<https://blog.qilahk.com/>

Reference editions and source pages checked October 12, 2026. The proposed mechanisms and evaluation plans in this paper are research designs, not reported experimental findings.